ACO Comprehensive Exam Part I Fall 2023

Aug 18, 2023

1 Graduate Algorithms

Given a connected, undirected, weighted graph $G = (V, E)$, a *maximum spanning tree* is a spanning tree of G of maximum total weight.

(a) (2 points) Given a connected, undirected, weighted graph $G = (V, E)$, with all edge weights distinct, prove the following *cycle property*:

For any cycle $\mathcal{C}$ of G, the edge of minimum weight in $\mathcal{C}$ does not belong to the maximum spanning tree.

(b) (3 points) Define the capacity of a path as the weight of the minimum edge in it. A bottleneck path between two vertices u and v is a path of maximum capacity among all paths connecting them. Design an algorithm that takes a weighted graph and pair of vertices u, v as input and outputs a bottleneck path between u and v. Analyze the runtime of your algorithm.

(c) (3 points) A network consists of a directed graph with nonnegative integer *capacities* $c(e)$ for every edge, along with a source vertex s and a sink vertex t. Given a network $\{G = (V, E); s, t \in V; \{c(e)\}_{e\in E}\}$, a flow is a function $f : E \to \mathbb{R}$ such that

(i) $0 \leq f(e) \leq c(e)$, for all $e \in E$

(ii) $\sum_{u\in V} f(uv) = \sum_{w\in V} f(vw)$ for all $v \in V \setminus \{s, t\}$

The size of a flow is defined as $s(f) = \sum_{w\in V} f(sw)$ (i.e., the flow out of the source). Recall the classical approach to find a flow of maximum size: find an *augmenting path* in the residual network and update the flow along this path. Repeat this procedure until there is no such paths. Using your procedure from part (b) to find an augmenting path, i.e., using a bottleneck (s, t)-path in the residual graph as the augmenting path, show that if the current flow is f and f^* is any maximum flow, then the max capacity path has value at least $\frac{s(f^*)-s(f)}{m}$, where $m = |E|$.

(d) (2 points) Use part (c) to derive a polynomial-time algorithm to find a Maximum Flow in the given network.

2 Linear Inequalities

1. Given a polyhedron $P \subseteq \mathbb{R}^n$, let

$$P' = \left\{x \in \mathbb{R}^n : c^T x \leq \lfloor \delta \rfloor \ \forall (c, \delta) \in \mathbb{Z}^n \times \mathbb{R} \text{ s.t. } c^T x \leq \delta \text{ is valid for } P\right\}$$

denote the Chvátal-Gomory closure of P and let $P_I = conv(P \cap \mathbb{Z}^n)$. For any integer $k \geq 0$, we also let $P^{(k)}$ denote the rank-k C-G-closure of P, defined as $P^{(0)} = P$, $P^{(1)} = P'$ and $P^{(k)} = \left(P^{(k-1)}\right)'$ for any $k \geq 2$. The C-G rank of P is the minimum integer k such that $P^{(k)} = P_I$.

Given a graph $G = (V, E)$, let

$$Q(G) = conv(\{\chi(F) : (V, F) \text{ is a connected spanning subgraph of } G\})$$

be the convex hull of connected spanning subgraphs where $\chi(F) \in \{0, 1\}^E$ is the indicator vector for edge set F. You may use the fact that

$$Q(G) = \{x \in R^E : x(\delta_G(\pi)) \geq |\pi| - 1 \ \forall \text{ partitions } \pi \text{ of } V, 0 \leq x \leq 1\}$$

where $\delta_G(\pi)$ denotes the set of edges with endpoints in different parts of the partition π and $|\pi|$ denotes the number of parts of π. Also, define the the cut-formulation

$$R(G) = \{x \in \mathbb{R}_+^E : x(\delta(S)) \geq 1 \ \forall \ \emptyset \subsetneq S \subsetneq V, 0 \leq x \leq 1\}$$

where $\delta(S)$ denotes the set of edges with exactly one endpoint in S.

(a) (1 point) Show that $R(G)_I = Q(G)$.

(b) (2 points) Show that there exist graphs where C-G rank of $R(G)$ is at least 1.

(c) (6 points) Show that for every integer k, $R(G)^{(k)}$ satisfies all inequalities of $Q(G)$ for partitions of size up to $k + 2$, i.e., show that

$$R(G)^{(k)} \subseteq \{x \in R^E : x(\delta_G(\pi)) \geq |\pi| - 1 \ \forall \text{ partitions } \pi \text{ of } V \text{ s.t.} |\pi| \leq k+2, 0 \leq x \leq 1\}$$

(d) (1 point) What does the above part imply about the C-G rank of $R(G)$.

3 Solutions

Algorithms.

(a) Let T be the maximum spanning tree (since all weights are distinct, it is guaranteed to be unique). Let $e = (uv)$ be the edge of minimum weight in a cycle $\mathcal{C}$. Assume $e \in T$. Consider the two connected components obtained by removing e from T. Since u, v will fall into different components and $\mathcal{C}$ is a cycle, there is at least one edge e' in $\mathcal{C}$ different from e that connects the two components (i.e., it is an edge in the *cut*). Now $T' = T \setminus \{e\} \cup \{e'\}$ is also a spanning tree and

$$w(T') = w(T) - w(e) + w(e') > w(T)$$

which contradicts the fact that T is maximum. Hence, our assumption is false and e is not in T.

(b) From part (a), we can deduce that *every path in the maximum spanning tree is a bottleneck path.* Indeed, let $\mathcal{P}$ be the bottleneck path connecting vertices u and v, and $\mathcal{P}_{uv}$ be the path connecting u and v in the maximum spanning tree. Assume $cap(\mathcal{P}) > cap(\mathcal{P}_{uv})$. This means the edge of minimum weight in $\mathcal{P}_{uv}$ is the edge of minimum weight in the union $\mathcal{P} \cup \mathcal{P}_{uv}$ (where the union is taking as a set of edges). But this in turns means the edge of minimum weight in the cycle resulting from the union above is in the maximum spanning tree, which contradicts (a). Now run Kruskal's algorithm sorting the edges in descending order and output the (unique) path connecting the given pair of vertices. The runtime is $O(|E|\log(|V|))$.

(c) Set $\alpha = \frac{s(f^*)-s(f)}{m}$. To see there is a path of capacity at least α in the residual network, drop all the edges of weight less than α from the residual network. If there is still a path connecting s and t, its capacity would be at least α, and we would be done. Otherwise, we have a cut that separates s and t and the remaining flow would be strictly less than $m\alpha = s(f^*) - s(f)$, contradicting the optimality of f^*.

(d) From (c), applying the above step, the remaining flow is at most

$$s(f^*) - s(f) - \frac{s(f^*) - s(f)}{m} = (1 - \frac{1}{m})(s(f^*) - s(f))$$

Thus, after ℓ rounds the flow left is at most $(1 - \frac{1}{m})^\ell s(f^*) < 1$ for $\ell = c\log(s(f^*))$. Finding the augmenting path takes $O(|E|\log(|V|))$ so our overall runtime is bounded by $O(|E|\log(|V|)\log(s(f^*))$, which is polynomial in the input size.

Linear Inequalities.

1. Every $x \in R(G) \cap \mathbb{Z}^E$ must be connected and therefore in $Q(G)$. Thus $R(G)_I \subseteq Q(G)$. Now consider any $\chi(F)$ where (V, F) is a spanning connected subgraph. Then clearly, $\chi(F) \in R(G)$ since it satisfies the cut constraint. Thus $\chi(F) \in R(G) \cap \mathbb{Z}^E$ and therefore $Q(G) \subseteq R(G)_I$ proving the equality.

2. Observe that C-G rank of $R(G)$ is zero iff $R(G) = R(G)_I = Q(G)$. Thus it is enough to find a graph where $Q(G) \subsetneq R(G)$ to show that C-G rank of $R(G)$ is at least one. Consider G to be a 3-vertex cycle and $x_e = \frac{1}{2}$ for each edge e of the cycle. Then $x \in R(G) \setminus Q(G)$.

3. We show, inductively, that the inequalities,

$$x(\delta_G(\pi)) \geq |\pi| - 1 \;\forall \text{ partitions } \pi \text{ of } V$$

are valid for $R(G)^{(k)}$ if $|\pi| \leq k + 2$. Clearly, the statement is true for $k = 0$. Consider any partition $|\pi| = k + 2$ where $k \geq 1$. Let $\pi = \{P_1, \ldots, P_{k+2}\}$. Consider the partition $\pi_{i,j}$ where we modify π by replacing P_i and P_j by a single part, the union of $P_i \cup P_j$, for each $1 \leq i < j \leq k$. Since, $\pi_{i,j}$ has $k + 1$ parts, by induction hypothesis we have that

$$x(\delta_G(\pi_{i,j})) \geq |\pi_{i,j}| - 1 = k$$

is valid for $R(G)^{(k-1)}$. Adding all these $\binom{k+2}{2}$ inequalities, we get that the following inequality is valid for $R(G)^{(k-1)}$:

$$\sum_{1 \leq i < j \leq k+2} x(\delta_G(\pi_{i,j})) \geq \binom{k+2}{2} \cdot k.$$

But we have

$$\sum_{1 \leq i < j \leq k+2} x(\delta_G(\pi_{i,j})) = \left(\binom{k+2}{2} - 1\right) x(\delta_G(\pi))$$

since every edge $\{u, v\} \in \delta_G(\pi)$ appears in all but one $\delta_G(\pi_{i,j})$. Thus we have that

$$x(\delta_G(\pi)) \geq \frac{\binom{k+2}{2}}{\binom{k+2}{2} - 1} \cdot k$$

is valid for $R(G)^{(k-1)}$. But then

$$x(\delta_G(\pi)) \geq \left\lceil \frac{\binom{k+2}{2}}{\binom{k+2}{2} - 1} \cdot k \right\rceil = k + 1$$

is valid for $R(G)^{(k)}$ as claimed.

4. Since all the constraints have $|\pi| \leq n$, all constraints defining $Q(G) = R(G)_I$ are valid for $R(G)^{(n-2)}$. Thus the C-G rank of $R(G)$ is at most $n - 2$.

ACO Comprehensive Exam Part II Fall 2023

Aug 17, 2023

1 Design and Analysis of Algorithms

- (4 points) Let $g : \mathbb{R}^n \to \mathbb{R}_+$ be a nonnegative L-Lipschitz function whose support is a convex body K of Euclidean diameter D. Consider the following algorithm to estimate the integral of g over K:

  1. Sample k uniform random points from K and compute the average value of $g(x)$ over the sample: $\mu_g = (1/k) \sum_{i=1}^{k} g(x_i)$.
  2. Output $\mu_g \cdot \mathrm{Vol}(K)$

  What is the expected output of the algorithm, with a single random sample? Assuming access to oracles for sampling K uniformly, and for its volume, how many samples k suffice to get a multiplicative $(1 + \epsilon)$-error approximation to the desired integral $\int_K g(x)\, dx$ with probability at least 3/4?

- (4 points) Now consider the problem of sampling from the density proportional to $g(x) = e^{-f(x)}$ restricted to a convex body K. Suppose you are given a membership oracle for K, i.e., numbers r, R and a point x_0 s.t. $x_0 + rB_n \subseteq K \subseteq RB_n$, and an oracle for evaluating f at any point $x \in K$. Consider the following algorithm:

  1. Sample (x, t) with density proportional to e^{-t} restricted to the set $\{(x, t) : x \in K, f(x) \le t\}$
  2. Output x.

  Show how to use the ball walk to implement Step 1. Then show that the marginal density of x is proportional to $e^{-f(x)}\chi_K(x)$. Give an interpretation of this algorithm as a reduction. You are not required to provide a convergence guarantee in this part.

- (2 points) What is the complexity status of sampling from the density proportional to $e^{-f(x)}$ restricted to K when f is convex, given by a function oracle and $\max_K f(x) - \min_K f(x) \le \beta$; K is convex, given by an (r, R)-membership oracle and a point x_0 s.t. $x_0 + rB_n \subseteq K \subseteq RB_n$ as above; and no other assumptions?

2 Combinatorial Optimization

Let $G = (V, E)$ be a undirected graph. Suppose that $E(S) \leq 2|S| - 2$ for all $S \subseteq V$ where $E(S) := \{\{u, v\} \in E : u, v \in S\}$ is the set of edges with both endpoints in S. Moreover, suppose that for some integer $k \geq 2$, G contains a spanning tree T such that the degree in T of each vertex $v \in V$ is at most k. In the following, we will see how to find a spanning tree H, in *polynomial time*, such that degree of each vertex is at most $k + 2$.

1. (4 points) Show that there exists an orientation $D = (V, A)$ of $G = (V, E)$ such that each vertex $v \in V$ has in-degree at most two in D.

2. (2 points) For any set of edges $F \subseteq E$, we let $\vec{F}$ denote the corresponding oriented arcs in D. Show that if H is a spanning tree of G such that $\vec{H} \subseteq A$ and every vertex $v \in V$ has out-degree at most k in $\vec{H}$, then degree of each vertex $v \in V$ is at most $k+2$ in H.

3. (4 points) Show that one can find in polynomial time, a spanning tree H of G such that degree of each vertex $v \in V$ in H is at most $k + 2$. You may use part (b) and the fact that there is a polynomial-time algorithm for matroid intersection.

3 Probabilistic Combinatorics

Problem. Suppose we throw n balls into n bins independently and uniformly at random. Let X be the random variable equal to the number of bins that remain empty.

(a) (5 points) Prove that $\mathrm{Var}[X] \le \mathbb{E}[X]$.

(b) Prove one of the following upper bounds for all $\lambda > 0$:

$$\Pr\left[|X - \mathbb{E}[X]| > \lambda\sqrt{n}\right] \le \begin{cases} e^{-1}\lambda^{-2} & \text{(partial credit, 2 points),} \\ \\ 2e^{-\lambda^2/2} & \text{(full credit, 5 points).} \end{cases}$$

4 Solutions

Design and Analysis of Algorithms.

- For a single random sample x drawn uniformly from K, the expected output is precisely

$$\mathbb{E}_K(g(x)) \cdot \mathrm{Vol}(K) = \int_K g(x)\, dx.$$

  The variance of the estimator $Y = f(x)$ is

$$\mathbb{E}_K((g(x) - \mathbb{E}_K(g(x))^2) \leq \mathbb{E}_K((\max_K g(x) - \min_K g(x))^2) \leq L^2 \dot{D}^2.$$

  By Chebychev's inequality, it suffices to use $O(L^2D^2/\epsilon^2)$ points to get additive error ϵ. However, to get multiplicative error, the number of samples will depend on the ratio of the variance to the square of the expectation. Although the expectation is positive (since g is positive over the domain), it could be arbitrariliy small, and so the number of samples would be

$$O\left(\frac{L^2D^2}{\epsilon^2\mathbb{E}_K(g(x))^2}\right).$$

- To implement the ball walk, generate a random point (y, s) in the ball of radius δ centered at (x, t); then go to (y, s) with probability $\min\{1, e^{-s}/e^{-t}\}$. Now let (x, t) be sampled with density $Ce^{-t}\chi_K(x)\chi(f(x) \leq t)$. The marginal along x is then precisely

$$\int_{t=f(x)}^{\infty} Ce^{-t}\chi_K(x)\, dt = Ce^{-f(x)}\chi_K(x).$$

- The problem can solved using $\mathrm{poly}(n, \log(R/r), \log(\beta/\epsilon))$ calls to the function evaluation and membership oracles as well as additional arithmetic operations.

Combinatorial Optimization.

1. By the Nash-Williams theorem, $G = (V, E)$ can be covered by two forests, say F_1 and F_2. For each tree in each of the forests, pick a root and direct all edges away from the root. Since each vertex appears in at most two trees, it receives at most two in-edges. This gives the orientation $D = (V, A)$.

2. If T is a spanning tree of G, then we have

$$deg_T(v) = deg^{out}_{\vec{T}}(v) + deg^{in}_{\vec{T}}(v) \leq k + deg^{in}_{\vec{T}}(v) \leq k + deg^{in}_{A}(v) \leq k + 2$$

   where the first inequality follows from assumption, second from the fact $\vec{T} \subseteq A$ and last inequality follows since in $D = (V, A)$ each vertex has at most 2 in-edges.

3. Consider the two matroids $\mathcal{M}_1 = (A, \mathcal{I}_1)$ and $\mathcal{M}_2 = (A, \mathcal{I}_2)$ where $\mathcal{M}_1$ is the graphic matroid where we ignore the directions on the arcs. Let $\mathcal{I}_2 = \{B \subseteq A : deg^{out}_B(v) \leq k \; \forall \, v \in V\}$. Observe that $\mathcal{M}_2$ is a partition matroid since the out-edges incident at each vertex form a partition of the ground set. Observe that $\vec{T}$ is a common independent set of size $|V| - 1$ since it is independent in both matroids. Thus we can find a common independent set $\vec{H} \subseteq A$ of size $|V| - 1$ using matroid intersection. Let $H \subseteq E$ denote the edges obtained after removing directions on edges in $\vec{H}$. Clearly, $\vec{H} \in \mathcal{M}_1$ implies H is a spanning tree. Moreover, $deg^{out}_{\vec{H}}(v) \leq k$ for each $v \in V$ since $\vec{H} \in \mathcal{I}_2$. Thus, from part (b), we have $deg_H(v) \leq k + 2$ for each $v \in V$.

Probabilistic Combinatorics. Let us denote the balls by $b_1, \ldots, b_n$ and the bins by $B_1, \ldots, B_n$.

(a) Let X_j be the indicator random variable of the event $\{B_j \text{ is empty}\}$, so $X = \sum_{j=1}^n X_j$. The probability that the ball b_i is in the bin B_j is $1/n$. Since the balls are thrown independently,

$$\mathbb{E}[X_j] \,=\, \Pr[B_j \text{ is empty}] \,=\, (1-1/n)^n \,:=\, q.$$

The linearity of expectation yields

$$\mathbb{E}[X] \,=\, \sum_{j=1}^{n} \mathbb{E}[X_j] \,=\, \left(1-\frac{1}{n}\right)^n.$$

To bound the variance of X, we use the formula

$$\mathrm{Var}[X] \,=\, \sum_{j=1}^{n} \mathrm{Var}[X_j] \,+\, \sum_{j \neq j'} \mathrm{Cov}[X_j, X_{j'}].$$

Since each X_j is a Bernoulli random variable with mean q, we have

$$\mathrm{Var}[X_j] \,=\, q(1-q) \,\leq\, q \,=\, \left(1-\frac{1}{n}\right)^n.$$

Also, for $j \neq j'$, we have

$$\begin{aligned}\mathrm{Cov}[X_j, X_{j'}] \,&=\, \mathbb{E}[X_j X_{j'}] \,-\, \mathbb{E}[X_j]\mathbb{E}[X_{j'}] \,=\, \P[B_j \text{ and } B_{j'} \text{ are both empty}] \,-\, q^2 \\ &=\, \left(1-\frac{2}{n}\right)^n - \left(1-\frac{1}{n}\right)^{2n} \,=\, \left(1-\frac{2}{n}\right)^n - \left(1-\frac{2}{n}+\frac{1}{n^2}\right)^n \,\leq\, 0.\end{aligned}$$

Putting everything together,

$$\mathrm{Var}[X] \,=\, \sum_{j=1}^{n} \mathrm{Var}[X_j] \,+\, \sum_{j \neq j'} \underbrace{\mathrm{Cov}[X_j, X_{j'}]}_{\leq 0} \,\leq\, n\left(1-\frac{1}{n}\right)^n \,=\, \mathbb{E}[X].$$

(b) *Proof of the quadratic bound*: Using part (a) and the Chebyshev inequality, we write

$$\Pr\left[|X-\mathbb{E}[X]| \,>\, \lambda\sqrt{n}\right] \,\leq\, \frac{\mathrm{Var}[X]}{\lambda^2 n} \,\leq\, \frac{\mathbb{E}[X]}{\lambda^2 n} \,=\, \left(1-\frac{1}{n}\right)^n \lambda^{-2} \,\leq\, e^{-1}\lambda^{-2}.$$

Proof of the exponential bound: Let $Y_i \in [n]$ be such that the ball b_i belongs to the bin B_{Y_i}. Then $Y_1, \ldots, Y_n$ are mutually independent random elements, and the random variable X is a 1-Lipschitz function of the tuple $(Y_1, \ldots, Y_n)$, since changing the value of Y_i for some $i \in [n]$ (i.e., moving the ball b_i to a different bin) can affect X at most by 1. Therefore, by Azuma's inequality,

$$\Pr[|X-\mathbb{E}[X]| > \lambda\sqrt{n}] \,\leq\, 2e^{-\lambda^2/2}.$$